

# NAS Parallel Benchmarks, Multi-Zone Versions

Rob F. Van der Wijngaart\*, Haoqiang Jin  
NASA Advanced Supercomputing (NAS) Division  
NASA Ames Research Center, Moffett Field, CA 94035-1000  
{wijngaar,hjin}@nas.nasa.gov

NAS Technical Report NAS-03-XXX

June 2003

## Abstract

We describe an extension of the NAS Parallel Benchmarks (NPB) suite that involves solving the application benchmarks LU, BT and SP on collections of loosely coupled discretization meshes. The solutions on the meshes are updated independently, but after each time step they exchange boundary value information. This strategy, which is common among structured-mesh production flow solver codes in use at NASA Ames and elsewhere, provides relatively easily exploitable coarse-grain parallelism between meshes. Since the individual application benchmarks also allow fine-grain parallelism themselves, this NPB extension, named NPB Multi-Zone (NPB-MZ), is a good candidate for testing hybrid and multi-level parallelization tools and strategies.

## 1 Introduction

The NAS Parallel Benchmarks (NPB) [1] are well-known problems for testing the capabilities of parallel computers and parallelization tools. They exhibit mostly fine-grain exploitable parallelism and are almost all iterative, requiring multiple data exchanges between processes within each iteration. Implementations in MPI [2], Java [4], High Performance Fortran [5], and OpenMP [6] all take advantage of this fine-grain parallelism.

However, many important scientific problems feature several levels of parallelism, and this property is not reflected in NPB. To remedy this deficiency the NPB Multi-Zone (NPB-MZ) versions were created, which are described in this report. Problem sizes and verification values are given for benchmark classes S, W, A, B, C, and D.

## 2 General Approach

The application benchmarks Lower-Upper Symmetric Gauss-Seidel (LU), Scalar Penta-diagonal (SP), and Block Tri-diagonal (BT) solve discretized versions of the unsteady, compressible Navier-Stokes equations in three spatial dimensions. Each operates on a structured discretization mesh that is a logical cube. In realistic applications, however, a single such mesh is often not sufficient to describe a complex domain, and multiple meshes or *zones* are used to cover it. In the production code OVERFLOW [3] the flow equations are solved independently in each zone, and after each iteration the zones exchange boundary values with their immediate neighbors with which they overlap.

---

\*Employee of Computer Sciences Corporation

## 2.1 Multi-Zone Mesh Systems

We take the OVERFLOW [3] approach in creating the NPB Multi-Zone versions of LU, BT, and SP, i.e. LU-MZ, BT-MZ, and SP-MZ. In each a logically rectangular discretization mesh is divided into a two-dimensional horizontal tiling of three-dimensional zones of approximately the same aggregate size as the original NPB, as indicated in Figure 1.




Figure 1: Two-dimensional tiling of three-dimensional mesh (exploded view).

Within all zones the LU, BT, or SP problems are solved to advance the time-dependent solution, using exactly the same methods and constants (except for mesh spacing, see Section 2.3) as described in [1, 7]. The mesh spacings of all zones of a particular problem class are identical, and the overlap between neighboring zones is exactly one such spacing, so that discretization points in overlap regions coincide exactly.

## 2.2 Data Exchange Between Zones

Exchange of boundary values between zones takes place after each time step, which provides the fairly loose coupling of the otherwise independent solution processes within the zones. The data transfer is as follows. Solution values at points one mesh spacing away from each vertical zone face are copied to the coincident boundary points of the neighboring zone. Values at points on zone edges and vertices are not used in the discretization formulas, so these do not need to be copied. The problem is periodic in the two horizontal directions ( $x$  and  $y$ ), so donor point values at the extreme sides of the mesh system are copied to boundary points at the opposite ends of the system. This data transfer process is shown schematically in Figure 2 for a horizontal two-dimensional slice of a mesh system of two by two zones. We only show where data residing on the bottom left slice is copied. No data is copied in the third space dimension, since the problem is not periodic in that direction, nor is the overall mesh system partitioned in that direction. The bottom left slice also receives data from other slices, but that is not shown in the figure.




Figure 2: Data transfer among zones within horizontal slice of three-dimensional mesh system. Only copying of data from bottom left zone slice is depicted.

### 2.3 Aggregate Mesh Sizes

To avoid pathologically shaped zones after partitioning the overall mesh in the two horizontal directions, we change the aspect ratios of the meshes of the original NPB. The total number of points in all zones is kept approximately the same as in the original NPB, except for the smallest problem class (S). There we increase the number of points to make sure that each zone contains enough points to fit the discretization stencil. The overall mesh sizes are listed in Table 1. The mesh spacing (distance between mesh points) used in the initializations and discretizations equals the reciprocal of the average number of mesh cells within each zone in the pertinent coordinate direction. The number of cells of any zone in a certain coordinate direction equals the number of points in that direction minus one. Hence, if the total number of points and the number of zones in a certain coordinate direction are  $np$  and  $nz$ , respectively, then the mesh spacing in that direction equals  $nz/(np - 1)$ .

Class	Mesh dimensions		
	$x$	$y$	$z$
S	24	24	6
W	64	64	8
A	128	128	16
B	304	208	17
C	480	320	28
D	1632	1216	34

Table 1: Total number of mesh points for all NPB multi-zone problems

## 3 Individual Benchmarks

The general approach described above covers most aspects of the NPB multi-zone problems. In this section we describe the differences between the three application benchmarks. We note that the

selection of different NPB solvers for the new benchmarks is fairly arbitrary. The major difference between the three multi-zone problems lies in the way the zones are created out of the single overall mesh.

### 3.1 LU-MZ

For all problem classes the number of zones in each of the two horizontal dimensions equals four. The overall mesh is partitioned such that the zones are identical in size, which makes it relatively easy to balance the load of the parallelized application. The actual sizes of the zones are listed in Table 2.

Class	dir	# zones	size
S	<i>x</i>	4	6
	<i>y</i>	4	6
	<i>z</i>	1	6
W	<i>x</i>	4	16
	<i>y</i>	4	16
	<i>z</i>	1	8
A	<i>x</i>	4	32
	<i>y</i>	4	32
	<i>z</i>	1	16
B	<i>x</i>	4	76
	<i>y</i>	4	52
	<i>z</i>	1	17
C	<i>x</i>	4	120
	<i>y</i>	4	80
	<i>z</i>	1	28
D	<i>x</i>	4	408
	<i>y</i>	4	304
	<i>z</i>	1	34

Table 2: Zone sizes (in points) for LU-MZ

Class	dir	# zones	size
S	<i>x</i>	2	12
	<i>y</i>	2	12
	<i>z</i>	1	6
W	<i>x</i>	4	16
	<i>y</i>	4	16
	<i>z</i>	1	8
A	<i>x</i>	4	32
	<i>y</i>	4	32
	<i>z</i>	1	16
B	<i>x</i>	8	38
	<i>y</i>	8	26
	<i>z</i>	1	17
C	<i>x</i>	16	30
	<i>y</i>	16	20
	<i>z</i>	1	28
D	<i>x</i>	32	51
	<i>y</i>	32	38
	<i>z</i>	1	34

Table 3: Zone sizes (in points) for SP-MZ

### 3.2 SP-MZ

As in the case of LU-MZ, the overall mesh is partitioned such that the zones are identical in size. However, the number of zones in each of the two horizontal dimensions grows as the problem size grows. The actual sizes of the zones are listed in Table 3.

### 3.3 BT-MZ

The number of zones in this benchmark grows with the problem size in the same fashion as in SP-MZ. However, the overall mesh is now partitioned such that the sizes of the zones span a significant range. This is accomplished by increasing sizes of successive zones in a particular coordinate direction in a roughly geometric fashion. An example of the stretched tiling of the overall mesh is shown in Figure 3. Except for class S, the ratio of largest over smallest total zone size is approximately 20. This makes it harder to balance the load than for SP-MZ and LU-MZ if the implementation is to take advantage of multi-level parallelism. The actual sizes of the successive zones for all problem classes are listed in Table 4.




Figure 3: Example of uneven mesh tiling (horizontal cut through mesh system) for the BT-MZ benchmark.

## 4 Performance Reporting

The paper-and-pencil specification of the original NPB defined a performance result of the individual application benchmarks as the elapse time of a certain segment of the whole code. A quantity that can be derived from that elapse time is the number of millions of floating point operations per second (MF/s). This is necessarily an approximation, since the actual number of floating point instructions executed depends on compiler technology and the presence of specialized hardware (floating point

Class	dir	# zones	sizes
S	$x$	2	6, 18
	$y$	2	6, 18
	$z$	1	6
W	$x$	4	6, 11, 18, 29
	$y$	4	6, 11, 18, 29
	$z$	1	8
A	$x$	4	13, 21, 36, 58
	$y$	4	13, 21, 36, 58
	$z$	1	16
B	$x$	8	16, 20, 24, 30, 38, 47, 57, 72
	$y$	8	11, 13, 17, 21, 26, 31, 40, 49
	$z$	1	17
C	$x$	16	13, 14, 15, 18, 19, 21, 23, 26, 28, 31, 35, 38, 43, 47, 52, 57
	$y$	16	8, 10, 10, 12, 12, 14, 16, 17, 19, 21, 23, 26, 28, 31, 35, 38
	$z$	1	28
D	$x$	32	22, 23, 24, 25, 26, 28, 29, 31, 32, 34, 35, 37, 39, 41, 43, 45, 48, 49, 52, 55, 58, 60, 63, 67, 70, 73, 77, 81, 85, 88, 94, 93
	$y$	32	16, 17, 18, 19, 20, 20, 22, 23, 24, 25, 26, 28, 29, 30, 33, 33, 35, 37, 39, 41, 43, 45, 47, 50, 52, 54, 58, 60, 63, 66, 70, 73
	$z$	1	34

Table 4: Zone sizes (in points) for BT-MZ

multiply/add units, division and square root hardware). In NPB 2.3 [2] curve fits were used that best fit data gathered by hardware performance counters on a number of systems available at that time. We use the same curve fits for computing MF/s for each zone of the new benchmark suite. While they are not very accurate since they were derived using meshes with different sizes and aspect ratios, they provide a convenient means of determining the relative performance of different systems. Let  $n_x$ ,  $n_y$ , and  $n_z$  be the extents of a zone in the three respective coordinate directions,  $n_i$  is the number of timed solver iterations, and  $T$  is the elapse time of the measured code segments in the application benchmarks. The following formulas define our performance curve fits.

$$\begin{aligned} \text{MF}/s_{LU} &= [1984.8n_xn_y n_z - 1213.7(n_x + n_y + n_z)^2 + 9257.0(n_x + n_y + n_z) - 144010] \frac{n_i}{10^6 T} \\ \text{MF}/s_{SP} &= [881.17n_xn_y n_z - 520.43(n_x + n_y + n_z)^2 + 3828.2(n_x + n_y + n_z) - 19272] \frac{n_i}{10^6 T} \\ \text{MF}/s_{BT} &= [3478.8n_xn_y n_z - 1961.7(n_x + n_y + n_z)^2 + 9341.2(n_x + n_y + n_z)] \frac{n_i}{10^6 T} \end{aligned}$$

## References

- [1] D. Bailey, E. Barszcz, J. Barton, D. Browning, R. Carter, L. Dagum, R. Fatoohi, S. Fineberg, P. Frederickson, T. Lasinski, R. Schreiber, H. Simon, V. Venkatakrishnan, S. Weeratunga. *The NAS Parallel Benchmarks*. NAS Technical Report RNR-94-007, NASA Ames Research Center, Moffett Field, CA, 1994.
- [2] D.H. Bailey, T. Harris, W.C. Saphir, R.F. Van der Wijngaart, A.C. Woo, M. Yarrow. *The NAS Parallel Benchmarks 2.0*. NAS Technical Report NAS-95-020, NASA Ames Research Center, Moffett Field, CA, 1995.
- [3] M.J. Djomehri, R. Biswas, M. Potsdam, R.C. Strawn. *An Analysis of Performance Enhancement Techniques for Overset Grid Applications*. NAS Technical Report NAS-03-008, NASA Ames Research Center, Moffett Field, CA, 2003.
- [4] M. Frumkin, M. Schultz, H. Jin, J. Yan. *Implementation of the NAS Parallel Benchmarks in Java*. NAS Technical Report NAS-02-009, NASA Ames Research Center, Moffett Field, CA, 2002.
- [5] M. Frumkin, H. Jin, J. Yan. *Implementation of NAS Parallel Benchmarks in High Performance Fortran*. NAS Technical Report NAS-98-009, NASA Ames Research Center, Moffett Field, CA, 1998.
- [6] H. Jin, M. Frumkin, J. Yan. *The OpenMP Implementation of NAS Parallel Benchmarks and its Performance*. NAS Technical Report NAS-99-011, NASA Ames Research Center, Moffett Field, CA, 1999.
- [7] R.F. Van Der Wijngaart. *NAS Parallel Benchmarks, Version 2.4*. NAS Technical Report NAS-02-007, NASA Ames Research Center, Moffett Field, CA, 2002.

## Appendix: Verification

The solution on each zone for each of the modified NPB problems is initialized in the same way as the single-zone NPB. Verification values are computed individually for each zone in the same way as for the original NPB, and are then summed over all zones, taking into account that  $L^2$  norms of solution errors and residuals are scaled using the actual number of points in each zone, not the overall number of points. The same error tolerances for verification apply as for the single-zone NPB.

Class	$m$	Residual norm	Error norm	surface integral
S	1	$0.5298535659341 * 10^1$	$0.2151113009059 * 10^2$	$0.4939825441335 * 10^2$
	2	$0.1197774899210 * 10^1$	$0.8618694694330 * 10^1$	
	3	$0.1994046926588 * 10^1$	$0.8304555534744 * 10^1$	
	4	$0.1075110730012 * 10^1$	$0.6965472514659 * 10^1$	
	5	$0.1187221359144 * 10^2$	$0.1106543938952 * 10^2$	
W	1	$0.8119553645098 * 10^3$	$0.6891894569748 * 10^2$	$0.3784788577400 * 10^3$
	2	$0.4418068636295 * 10^2$	$0.1330618340364 * 10^2$	
	3	$0.1179300513887 * 10^3$	$0.1095236968626 * 10^2$	
	4	$0.1561307465561 * 10^3$	$0.1323487700455 * 10^2$	
	5	$0.1683266367841 * 10^4$	$0.1781981334164 * 10^3$	
A	1	$0.1151708655328 * 10^4$	$0.1087735464068 * 10^3$	$0.6032448785219 * 10^3$
	2	$0.8597477575549 * 10^2$	$0.1870682342245 * 10^2$	
	3	$0.1452342063773 * 10^3$	$0.1403300551676 * 10^2$	
	4	$0.2144920032008 * 10^3$	$0.1948601652005 * 10^3$	
	5	$0.2254265852581 * 10^4$	$0.2481179295370 * 10^3$	
B	1	$0.2069632608421 * 10^5$	$0.1844564463808 * 10^4$	$0.5961640408886 * 10^4$
	2	$0.1012033769055 * 10^4$	$0.2987979377037 * 10^2$	
	3	$0.4279269356550 * 10^4$	$0.3197409251178 * 10^3$	
	4	$0.4058284641847 * 10^4$	$0.3606221167656 * 10^3$	
	5	$0.3850410135032 * 10^5$	$0.3646331526559 * 10^4$	
C	1	$0.5028123127554 * 10^5$	$0.3782792942929 * 10^4$	$0.1102699916987 * 10^5$
	2	$0.4556296757509 * 10^4$	$0.1384589914538 * 10^3$	
	3	$0.1209287096203 * 10^5$	$0.7726355958805 * 10^3$	
	4	$0.9611757784900 * 10^4$	$0.7576405186756 * 10^3$	
	5	$0.8459506446641 * 10^5$	$0.7332086531811 * 10^4$	
D	1	$0.2867560282080 * 10^6$	$0.7178690396542 * 10^4$	$0.2033260579515 * 10^5$
	2	$0.3426906735053 * 10^5$	$0.5393515916385 * 10^3$	
	3	$0.8112265314207 * 10^5$	$0.1695009242039 * 10^4$	
	4	$0.6443400298538 * 10^5$	$0.1479308943254 * 10^4$	
	5	$0.3957537402382 * 10^6$	$0.1274675411762 * 10^5$	

Table 5: Verification values for LU-MZ.  $m$  signifies vector component.

Class	$m$	Residual norm	Error norm
S	1	$0.8793040366889 * 10^1$	$0.7389283698609 * 10^1$
	2	$0.1784013414719 * 10^1$	$0.3391062571875 * 10^0$
	3	$0.2957988275205 * 10^1$	$0.3284776719979 * 10^1$
	4	$0.2395846121792 * 10^1$	$0.2707014619466 * 10^1$
	5	$0.1605964581201 * 10^2$	$0.5256684316018 * 10^1$
W	1	$0.1850663146585 * 10^3$	$0.1770251933254 * 10^2$
	2	$0.7403086025540 * 10^1$	$0.1258429151324 * 10^2$
	3	$0.3043754049324 * 10^2$	$0.1068623656417 * 10^2$
	4	$0.3362127636254 * 10^2$	$0.8840693457674 * 10^1$
	5	$0.3883549257008 * 10^3$	$0.5084604992926 * 10^2$
A	1	$0.2758386354855 * 10^3$	$0.1966639858342 * 10^2$
	2	$0.1604873427686 * 10^2$	$0.1184816849622 * 10^2$
	3	$0.4229764253213 * 10^2$	$0.9170443497284 * 10^1$
	4	$0.5088027003753 * 10^2$	$0.7345834613203 * 10^1$
	5	$0.5419995013086 * 10^3$	$0.6796481878493 * 10^2$
B	1	$0.5035539171425 * 10^4$	$0.5128417836351 * 10^3$
	2	$0.2760465367064 * 10^3$	$0.7164181156822 * 10^2$
	3	$0.7668837008516 * 10^3$	$0.7458453168794 * 10^2$
	4	$0.9776142499181 * 10^3$	$0.9499787054727 * 10^2$
	5	$0.1036235395807 * 10^5$	$0.1206268286372 * 10^4$
C	1	$0.5765173908781 * 10^5$	$0.6163567187924 * 10^4$
	2	$0.1897550071826 * 10^4$	$0.4231966935222 * 10^3$
	3	$0.9677613124492 * 10^4$	$0.9497955167739 * 10^3$
	4	$0.1181354420896 * 10^5$	$0.1200340600161 * 10^4$
	5	$0.1236659963809 * 10^6$	$0.1382265435960 * 10^5$
D	1	$0.8082004759829 * 10^6$	$0.8105627621632 * 10^5$
	2	$0.1591006213080 * 10^5$	$0.2110642763606 * 10^4$
	3	$0.1568266208093 * 10^6$	$0.1371409543331 * 10^5$
	4	$0.1673622751981 * 10^6$	$0.1629181880012 * 10^5$
	5	$0.1660834041393 * 10^7$	$0.1716124395911 * 10^6$

Table 6: Verification values for SP-MZ.  $m$  signifies vector component.



Class	$m$	Residual norm	Error norm
S	1	$0.1011775452720 * 10^4$	$0.1724390108088 * 10^3$
	2	$0.3650608110905 * 10^2$	$0.1041542956290 * 10^2$
	3	$0.1511220409149 * 10^3$	$0.2888900785025 * 10^2$
	4	$0.1370836311993 * 10^3$	$0.2539696848540 * 10^2$
	5	$0.1093080581272 * 10^4$	$0.1915299420916 * 10^3$
W	1	$0.4615016845424 * 10^5$	$0.5697492736075 * 10^4$
	2	$0.2682488801830 * 10^4$	$0.2862381824415 * 10^3$
	3	$0.7632320101012 * 10^4$	$0.9157387079752 * 10^3$
	4	$0.5726358936658 * 10^4$	$0.7150871042023 * 10^3$
	5	$0.4119836325082 * 10^5$	$0.5256642020838 * 10^4$
A	1	$0.5037047616879 * 10^5$	$0.5687733820812 * 10^4$
	2	$0.3033355278710 * 10^4$	$0.2790908558747 * 10^3$
	3	$0.8271031115772 * 10^4$	$0.8915884422633 * 10^3$
	4	$0.6041186307729 * 10^4$	$0.6906145002891 * 10^3$
	5	$0.4205478745176 * 10^5$	$0.4945592388790 * 10^4$
B	1	$0.4843459102489 * 10^6$	$0.4626290731551 * 10^5$
	2	$0.3818923043701 * 10^5$	$0.2759954836339 * 10^4$
	3	$0.8698756134349 * 10^5$	$0.7781825678461 * 10^4$
	4	$0.5841702114834 * 10^5$	$0.5810069914308 * 10^4$
	5	$0.4022899173953 * 10^6$	$0.4265661900427 * 10^5$
C	1	$0.4237283258604 * 10^7$	$0.2273994376606 * 10^4$
	2	$0.3295341366857 * 10^6$	$0.1549759512506 * 10^5$
	3	$0.7061306845572 * 10^6$	$0.4072489322623 * 10^5$
	4	$0.4523474014212 * 10^6$	$0.3028564316114 * 10^5$
	5	$0.2706046124276 * 10^7$	$0.2276701656056 * 10^6$
D	1	$0.4371642606971 * 10^8$	$0.1021827173069 * 10^7$
	2	$0.4151827778194 * 10^7$	$0.8116900108877 * 10^5$
	3	$0.8678190487727 * 10^7$	$0.1983571722729 * 10^6$
	4	$0.5785889993558 * 10^7$	$0.1438497054698 * 10^6$
	5	$0.3090315869859 * 10^8$	$0.1030366887428 * 10^7$

Table 7: Verification values for BT-MZ.  $m$  signifies vector component.